

Attorney Docket: TN112

REMARKS

Claims 1-11 are currently pending in the subject application. More specifically, claims 1-3 and 5-11 are rejected under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,287,453 (hereinafter, "Roberts"). Claim 4 is rejected under 35 U.S.C. §103(a) as being unpatentable over Roberts. Although Applicant respectfully disagrees with the Examiner's assertions set forth in the *Response to Arguments* section of the Office Action dated March 4, 2004, (e.g., administrative control within each node is *not* solely relevant to MSCS), Applicant has amended independent claims 1, 5, and 8 herein in order to more clearly delineate the intended scope of his claimed invention.

I. Background

A. Applicant's Invention

Applicant has previously set forth a summary of the claimed invention of the subject application, but it is restated here for the ease of the Examiner's reference.

Applicant's claimed invention seeks to provide a solution where a distributed application that requires centralized administration via a master node is running in a clustered computing environment, where administrative control resides in each node of the environment. More specifically, as set forth in detail in the subject specification, it is appreciated that distributed network applications often have defined the concept of a master machine that performs administration for the entire distributed application. (*Application No. 09/127,167*, p. 2) In the exemplary case of the Tuxedo environment, one of these Logical Machines is designated as the master, on which is running a DBBL process which performs administration for the entire Domain, including bringing a component online, taking a component offline, or checking the status of an individual component. (*Application No. 09/127,167*, p.3).

A problem arises with such a distributed application when it is required to run in a clustered computing environment, such as by way of example, Microsoft Cluster Server (MSCS) in Microsoft® Windows NT®, Enterprise Edition, where administration is implemented on each of the connected systems (nodes) composing the cluster. As set forth in more detail in the specification, in the exemplary clustered environment, MSCS is controlled by the Cluster Service, which runs on each note of the cluster. (*Id.*). The Cluster Service spawns one or more Resource Monitors, each of which calls entry points in a Resource DLL, the latter of which implements the actions needed to bring the resource on-line or to take the resources off-line. (*Id.*). Thus, contrary to distributed network application that requires centralized

Attorney Docket: TN112

administration via a master node, each node in the clustered computing environment maintains the administrative control.

Applicant's invention provides a solution to this problem. As set forth in claims 1, 5, 8, and 9, an administrative request from the clustered computing environment is first received at an originating node, and it is then determined whether the originating node is a designated master node for the distributed network application. If it is determined that the originating mode is not the designated master node, then the administrative request is routed to that mode which is the designated master node. (See, *Serial No. 09/127,167, claims 1, 5, 8 and 9*). In a preferred embodiment, this is accomplished an instance of a named pipe is created between the originating node and the designated master node, and passing the administrative request from the originating node to the designated master node via the named pipe. (See, *Serial No. 09/127,167; e.g., claims 2, 6, and 10*).

B. Newly Cited Art – Roberts

Roberts discloses a fast remote file access facility for distributing file access requests in a closely-coupled distributed system that includes a plurality of independently-operated computer systems located in close proximity to each other. (*Roberts, Abstract; Cols. 1-2, lines 65-68 and 1-3*). More specifically, the system of Roberts shows a closely-coupled computer system including a cluster driver and a high speed cluster controller. The driver provides an interface to the cluster controller that is operatively connected to a fast remote file access (FRFA) facility which is included as part of the operating system of each system when the system is enabled to operate in cluster mode. (*Roberts, Col. 2, lines 9-17*). The FRFA facility responds to various monitor or system calls from user application programs executing on a system. If it is determined that the user requires the use of a peripheral device attached to another computer system, the FRFA facility uses the existing file management facilities of the operating system to generate a message designating the complete path and specific peripheral device in the other system. This message is then transferred through the cluster driver and cluster controller to the FRFA facility of the other computer system. This operation is accomplished in a manner which makes the sets of peripheral devices of the other computer systems similarly enabled to operate in cluster mode appear to the user to be locally attached to the system bus of that system. (*Id. at lines 19-34*) According to Roberts, one of the primary advantages of his invention is that it "eliminates the need to continually issue messages since the determination of whether the object is a local or remote peripheral device is carried out at a sufficiently high level so as to eliminate the need for additional analysis." (*Id. at lines 35-39*)

Attorney Docket: TN112

II. Rejection of Claims 1-3, and 5-11 under 35 U.S.C. §102(b) - (Roberts)

A. Roberts Does Not Teach or Disclose Appellant's Claimed Invention of Enabling a Distributed Network Application that Requires Central Administration Via a Defined Master Node to Execute on the Nodes of a Clustered Computing Environment of the Type Wherein Each Node Generally Implements its Own Administrative Operations

Although Roberts' invention involves closely-coupled distributed systems, the invention is not concerned with enabling a distributed network application that requires central administration via a defined master node to execute on the nodes of a clustered computing environment of the type wherein each node generally implements its own administrative operations. There is no teaching or suggestion of distributed applications and the central administrative functions performed thereunder in Roberts. As there is no teaching or suggestion of this aspect of Applicant's claimed invention, claims 1-11 are patentable over Roberts.

1). The Automatic Volume Recognition (AVR) Software Cited by the Examiner Is Not a Distributed Application as is Generally Defined by Those of Skill in the Art and as Used by Appellant in the Subject Application

The Examiner indicates at page 2, however, that the automatic volume recognition (AVR) software discussed in Roberts is equivalent to a "distributed application." (see, e.g., *Office Action* dated March 4, 2004, p. 2, ¶4). Applicant respectfully disagrees with the Examiner's categorization of the AVR software, as it is clear from Roberts that the software is not "distributed" as is commonly used in the art and as used by Applicant in the subject application.

As known to those skilled in the art, and as used in the subject patent application, a "distributed application" is a program that generally comprises two or more parts and is designed to run on more than one computer in a network. (see, e.g., <http://e-docs.bea.com/wle/tuxedo/glossary/glossary.htm>, http://java_sun.com/docs/books/tutorial/idl/summary/_jidlGlossary.html, and http://www.liebert.com/support/glossary/net_gloss.asp copies of relevant pages being attached hereto as Appendix A).

As stated above, each host processor in Roberts includes a driver and fast remote file access (FRFA) facility. (See, e.g., Figures 1 and 3 and accompanying description). Each FRFA facility on each processor includes, *inter alia*, a file system manager module 12-240 on which is executing automatic volume recognition (AVR) software, which operates to recognize when a disk volume has been mounted or dismounted locally. (*Roberts*, Col. 6, lines 6-9). According to

Attorney Docket TN112

the invention of Roberts, the members of a cluster exchange a list of all of the disks, tapes and printers to be shared across the cluster, which eliminates the need for having to build a catalog identifying all resources to be referenced by a local member. The AVR software recognizes the volume locally and sends a message to all other cluster members informing them that the named volume has been mounted at that member. Assuming that it is the only disk pack so named, it will be identified as a local volume at each member in the cluster upon the issuance of a device status (STS) command. (*Roberts*, Col. 7, lines 15-19 and 22-26). Importantly, each processor includes a system manager module 12-240 on which is executing its own automatic volume recognition (AVR) software:

*When a local volume/device has been dismounted the AVR software of a cluster member sends a message to its CMGR task module (*Roberts*, Col. 10, lines 15-17)(Emphasis added)*

Therefore, as each processor includes a system manager module 12-240 on which is executing its own automatic volume recognition (AVR) software, the AVR software cannot be said to be equivalent to a "distributed application" as that term is generally understood, and as used by Applicant in the subject patent application. Thus, it will be appreciated Roberts does not teach enabling a distributed network application that requires central administration via a defined master node to execute on the nodes of a clustered computing environment of the type wherein each node generally implements its own administrative operations. For this reason alone, claims 1-11 are patentable over Roberts.

2). Roberts Teaches Away From Applicant's Claimed Invention Since A Master Node for a Distributed Application in Applicant's Claimed Invention is Defined and Distinct and Does Not Change with Active State of a Cluster Member

In addition to the foregoing reasons, claims 1-11 are patentable over Roberts because Roberts teaches away from Applicant's claimed invention of defined, distinct master node that performs administration for the entire distributed application. Instead, Roberts teaches a "master member" that is not defined/distinct and changes with the active state of a cluster member. More specifically:

Additionally, CMGR module 12-220 includes a routine for performing an I-AM_master function. This function is executed after an existing cluster master member fails and a new cluster master member must be determined. It determines whether or not the cluster member is the "master" within the system. This is determined by calculating the number of the member which has the lowest member number in the cluster (first bit in an active or on state in a bit map). This number is then compared to the local member number. If the local member number is less than the lowest number in the bit map, then the local member is the master in which case the routine returns a true or ONE value. For example, if cluster member numbers 2, 4 and 6 are active and member 6 fails or is

Attorney Docket: TN112

shut down, then member 2 becomes the new master. (Roberts, Col. 9, lines 30-46)(Emphasis added)

Thus, unlike Applicant's defined, distinct master node, it can be seen in Roberts that the a cluster member may be redesignated as "master" in at least those circumstances where an existing cluster master member fails. For this additional reason, claims 1-11 are patentable over Roberts.

B. Roberts Does Not Teach or Suggest the Step of Determining for Each Received Administrative Request Whether the Originating Node is a Designated Master Node for the Distributed Application and then Routing This Administrative Request from the Originating Node to the Designated Master Node if the Originating Node is not the Designated Master Node

In addition to the foregoing reasons, Roberts' invention further does not teach or suggest the step of determining for each received administrative request whether the originating node is a designated master node for the distributed application and then routing this administrative request from the originating node to the designated master node if the originating node is not the designated master node. As there is no teaching or suggestion of this aspect of Applicant's claimed invention, claims 1-11 are patentable over Roberts.

1). Although Applicant's Determining and Routing Steps are Clearly Directed to the Same Administrative Request, the Examiner Improperly Chooses Two Separate and Distinct Administrative Requests in Roberts to Provide Support for the Rejection.

The Examiner improperly cites two separate and unrelated administrative requests from different modules in the Roberts reference to provide evidence of teaching the determining and routing steps of independent claims 1, 5, and 8 of the subject application (*see, generally, Office Action* dated March 4, 2004). More specifically, in providing support for the teaching of Applicant's step of determining for each received administrative request whether the originating node is a designated master node for the distributed application, the Examiner cites Roberts at Col. 10, lines 30-67:

The CMON module 12-230 includes routines for performing a number of functions which primarily involve monitoring the state of cluster controller 14, through cluster controller driver 12-300 of FIG. 2 and communicating any changes. These routines include a get.sub.--cluster.sub.--lrn function wherein a routine returns the logical resource number (LRN) associated with the cluster controller 14 identified as "ICLUS00", a disconnect function wherein another routine enables the CMON module 12-230 to disconnect from the cluster controller 14 so that it can reconnect when the controller 14 becomes available and a connect.sub.--to function wherein the CMON module 12-230 is allowed to connect to the cluster controller 14. Once connected, CMON module 12-230 checks a status message received from cluster controller 14 to see if it is the master and whether there are other members already connected to the controller 14. If the status message indicates that it is the only active member, then it assumes it is the master. The

Attorney Docket: TN112

CMON module sends a message to the CMGR module indicating whether or not it is the master. (Roberts, Col. 10, lines 31-51)(Emphasis added)

With regard to providing evidence of teaching of the step of routing this administrative request from the originating node to the designated master node if the originating node is not the designated master node, the Examiner then turns to Cols. 9-10, lines 57-68; lines 1-30, stating that this limitation is met by Roberts' teaching that a message is sent to the CMGR module 12-220 of the master member to determine if the volume can be mounted cluster wide (*Office Action* dated March 4, 2004, p. 3):

An AVR.sub.--action routine also included as part of the CMGR module is used whenever a message is received from the AVR software. When a message is received indicating that a volume has been mounted, it is the responsibility of the CMGR module to check with the master (if it is not the master) to determine if the volume can be mounted cluster wide (has a name with a unique value). This involves sending a message to the CMGR module of the master member via control pipe 12-250 SCATALOG>CM.PIPE. This results in the CMGR module of the master member sending a message to the CMGR module of each cluster member to determine if the name of the mounted volume/device is unique across the cluster. When the CMGR module of the master member determines that the mounted volume/device is unique, it then sends an appropriate message to the CMGR module of each cluster member indicating its approval of the mounted volume/device. This results in each CMGR module issuing an ACTIVATE volume call to its monitor call handler module 12-140. In response to such call, the cluster access section of file manager module 12-240 of each cluster member creates a RDB structure for the mounted volume/device.

When a local volume/device has been dismounted the AVR software of a cluster member sends a message to its CMGR task module. In a similar to that manner described above, the CMGR module communicates with the master member which in turn informs the other members of the cluster to delete the RDB structure representative of the dismounted volume/device.

The CMGR module 12-220 keeps track of all of the semimounted volumes on its system that have not been approved for mounting by the master in the event of a master failure. Lastly, the CMGR module 12-220 includes routines for performing cluster, member and device access operations (Roberts, Cols. 9-10, lines 57-68; lines 1-30)

Clearly, the these administrative requests – the CMON module 12-230 checking to see if it is a master, and the CMGR module 12-220 checking with the master member (if it is not the master) to determine if a volume can be mounted cluster wide – are distinct and unrelated, administrative requests. Conversely, in Applicant's claimed invention, the determining and routing steps involve the same administrative request. For this additional reason, claims 1-11 are patentable over Roberts.

Attorney Docket: TN112

2). Roberts Does Not Teach or Suggest the Step of Determining for Each Received Administrative Request Whether the Originating Node is a Designated Master Node for the Distributed Application.

Even were one to ignore the arguments set forth in II.B.1. above, Roberts still does not teach Applicant's claimed limitations of determining for each received administrative request whether the originating node is a designated master node for the distributed application and then routing this administrative request from the originating node to the designated master node if the originating node is not the designated master node.

Again, the Examiner states that Applicant's claimed limitation of the step of determining whether the originating node is a designated master node for the distributed application is met by Roberts' teaching that the CMON module 12-230 sends a message to the CMGR module 12-220 indicating whether it is a master. (*See, Office Action dated March 4, 2004; see also, Roberts, Col. 10, lines 31-51*). However, Applicant respectfully disagrees with the Examiner. Contrary to Applicant's claimed invention, Roberts does not teach or disclose Applicant's step of determining for each received administrative request whether the originating node is a designated master node for the distributed application.

Referring again to the relevant section of Roberts:

The CMON module 12-230 includes routines for performing a number of functions which primarily involve monitoring the state of cluster controller 14, through cluster controller driver 12-300 of FIG. 2 and communicating any changes. These routines include a get.sub.--cluster.sub.--lrn function wherein a routine returns the logical resource number (LRN) associated with the cluster controller 14 identified as "/CLUS00", a disconnect function wherein another routine enables the CMON module 12-230 to disconnect from the cluster controller 14 so that it can reconnect when the controller 14 becomes available and a connect.sub.-- to function wherein the CMON module 12-230 is allowed to connect to the cluster controller 14. Once connected, CMON module 12-230 checks a status message received from cluster controller 14 to see if it is the master and whether there are other members already connected to the controller 14. If the status message indicates that it is the only active member, then it assumes it is the master. The CMON module sends a message to the CMGR module indicating whether or not it is the master. (Roberts, Col. 10, lines 31-51)(Emphasis added)

Thus, the only time when CMON module 12-230 checks with the cluster controller 14 to see if it is a master and then send this message to the CMGR module 12-220, occurs upon connection to the cluster of the relevant cluster member in which the CMON module 12-230 resides. Conversely, Applicant's claimed invention teaches that the step of whether the originating node is a designated master node for the distributed application is determined for each received administrative request. Although Applicant believes that this was clear in the original wording of the claims, he has amended the relevant

Attorney Docket: TN112

claims to more clearly indicate that this is the intent of the scope of the invention. As there is no teaching or suggestion of this aspect of Applicant's claimed invention, claims 1-11 are patentable over Roberts.

3). Roberts Does Not Teach or Suggest the Step of Routing Each Administrative Request from the Originating Node to the Designated Master Node if the Originating Node is not the Designated Master Node.

Roberts does not teach or suggest the step of routing each administrative request from the originating node to the designated master node if the originating node is not the designated master node. As set forth in the section above, the Examiner states that this step of Applicant's claimed invention is disclosed in Roberts' teaching that a message is sent to the CMGR module 12-220 of the master member to determine if the volume can be mounted cluster wide (*Office Action* dated March 4, 2004, p. 3). Even to determine if the volume can be mounted cluster wide (*Office Action* dated March 4, 2004, p. 3). Even ignoring the arguments set forth in II.B.1. above, Applicant respectfully disagrees for at least the following reasons.

First, as set forth above, Roberts' teaching is limited to *only* those messages sent to the CMGR module 12-220 of the master member and *only* in the instance when a volume is being loaded. As set forth in Applicant's claimed invention, the step of routing is applicable to *each* administrative request from the originating node to the designated master node if the originating node is not the designated master node.

In addition, in Applicant's claimed invention, if it is determined for an administrative request that the originating node is *not* the designated master node, the administrative request (itself) is routed from the originating node to the designated master node. Conversely, the receiving CMGR module 12-220 of Roberts does not route the message it receives from the AVR software to the master member CMGR module 12-220, but rather sends a separate message:

An AVR.sub.- action routine also included as part of the CMGR module is used whenever a message is received from the AVR software. When a message is received indicating that a volume has been mounted, it is the responsibility of the CMGR module to check with the master (if it is not the master) to determine if the volume can be mounted cluster wide (has a name with a unique value). This involves sending a message to the CMGR module of the master member via control pipe 12-250 SCATALOG>CM.PIPE. This results in the CMGR module of the master member sending a message to the CMGR module of each cluster member to determine if the name of the mounted volume/device is unique across the cluster. When the CMGR module of the master member determines that the mounted volume/device is unique, it then sends an appropriate message to the CMGR module of each cluster member indicating its approval of the mounted volume/device. This results in each CMGR module issuing an ACTIVATE volume call to its monitor call handler module 12-140. In response to such

Attorney Docket: TN112


call, the cluster access section of file manager module 12-240 of each cluster member creates a RDB structure for the mounted volume/device. (Roberts, Cols. 9-10, lines 59-68; 1-14)

Therefore, Roberts also does not teach Applicant's claimed limitation of a step of routing each administrative request from the originating node to the designated master node if the originating node is not the designated master node. For this additional reason, claims 1-11 are patentable over Roberts.

II. Conclusion

Applicant respectfully submits that claims 1-11 are patentable under 35 U.S.C. §102(b) and 35 U.S.C. §103(a) Roberts. Reconsideration of these claims and the application as a whole is thus respectfully solicited. It is believed that all of the stated grounds of rejection have been properly traversed, accommodated, or rendered moot. It is further believed that a full and complete response has been made to the outstanding Office Action, and as such, the present application, including claims 1-11, is in condition for allowance. Applicants therefore respectfully request prompt and favorable consideration of this amendment. If the Examiner believes that personal communication will expedite prosecution of this application, the Examiner is invited to telephone the undersigned at (215) 986-5169.

Respectfully submitted,




Lise A. Kode
Attorney for Applicant
Reg. No. 37,226

Unisys Corporation
Unisys Way, M/S E8-114
Blue Bell, Pennsylvania 19424-0001

The Director for Patents is hereby authorized to charge payment to Deposit Account No. 19-3790 of any fees associated with this communication.

I hereby certify that this correspondence is being transmitted via facsimile ((703) 872-9306) to the United States Patent and Trademark Office on the date shown below.

September 3, 2004



Serial No. 09/127,167

13